

The history of democracy and economics has proved that centralized power stifles human potential.

The history of democracy and economics has proved that centralized power stifles human potential.



I propose a philosophy based on individual empowerment as the source of prosperity, invention as the primary

purpose of superintelligence, and balance of power as the foundation of safety.

It is surprising that the discourse from many of those who are developing artificial intelligence is so filled with doom. I don't understand why anyone who believes that AI will eliminate most jobs and much of humanity's relevance would rush to build that future. The notion that AI is so dangerous that the only safe path is an extreme concentration of power seems dangerous. Historically, hoping that an absolute power will benevolently provide for humanity if sufficiently enlightened hasn't led to safe or positive outcomes.

Humanity has witnessed many transformative advances. Each time there is fear that people will be left behind. But each time humanity has come out with more people sharing greater prosperity, health and freedom. We believe this will be true with AI as well, and the abundance of the future can be shared by everyone. We also believe that the values that got us to this point—liberty, open inquiry, free enterprise and equal opportunity—are also the values to build a positive future.

Putting power in people's hands to pursue their own aspirations is how humanity has made the most progress. Novel ideas and major steps forward rarely originate from established institutions alone. They came from the brothers in a bicycle shop who believed people could fly, the bookbinder's apprentice with no schooling who figured out how to generate electricity, and the kid in a garage who thought personal computers could be for everyone. We believe this will continue to be true. As everyone gains more powerful tools, each person will become more capable of shaping the future, not less.

Invention, not automation, will be the greatest contribution of superintelligence. Early AI could answer questions and do routine work. Soon it will increasingly help discover new knowledge—from discovering new drugs to cure a family member's disease to finding new ways to improve your business. While the number of questions you can ask in a day is limited, the number of valuable things superintelligence can invent to help achieve your goals is unlimited.

As intelligence becomes abundant, the most important question will be how we direct it. Some argue that superintelligence itself, or a small set of experts who control it, should decide what is best for humanity. I disagree. The history of democracy and economics has shown that there is no single objective answer to how people define the best life, and therefore the best approach is letting people decide what matters to them.

Rather than centralizing this power, we believe that delivering personal superintelligence to everyone is the way to answer this question. This has the potential to begin a new era of personal empowerment, in which individuals have greater freedom to pursue their interests and reach their full potential.

Healthy societies also recognize that safety requires checks and balances. If only a handful of institutions have superintelligence, they will inevitably exercise a controlling influence over economics, science and politics. Even with the best intentions, that concentration would limit people's ability to choose their own future.

As a thought experiment, imagine only one person had a superintelligent lawyer. He would have an unfair advantage in court—even if his position was wrong on the merits. That would lead to a worse society. But now imagine everyone has a superintelligent lawyer. Justice would be carried out much more fairly and efficiently than it is today. When people are empowered, they naturally check and balance each other and the power of larger institutions.

Most societal questions can't be solved by attempting to align a single benevolent superintelligence. Humanity isn't a monoculture. There is no technological solution that can align with everyone's opposing interests and diverse values at once. Any singular superintelligence would have to prioritize some values over others and in the process would be incapable of being benevolent to everyone.

Many risks require serious attention. In most cases, like cybersecurity, the history of open-source software has shown that giving everyone full access to powerful systems will be the best way to protect safety and security over time. For other risks, including biological risks, more coordination between governments and other institutions on responsibly deploying capable models will be helpful.

This applies to the economy and the future of jobs as well. There is a balance between the use of AI as a tool for automation versus as a tool for empowering people to expand their skills and build new businesses. If the balance leans toward automation, the impact on jobs and the economy may be negative.

But if superintelligence is widely distributed, then I believe we will see more jobs in the future, not fewer. It will be significantly easier to start businesses without raising large amounts of capital. I expect the economy will become more entrepreneurial with a greater number of people working at small businesses rather than larger companies.

Developing superintelligence will be the most profound technological advance we will see in our lifetimes. Meta is committed to building with the principles of individual empowerment, invention and balance of power. The arc of human history has bent toward putting more power in people's hands. If these values lead the way, then I am optimistic that we can build a positive future for everyone.

Mr. Zuckerberg is CEO of Meta.